

WHAT IS CLAIMED IS:

1. A network switching device comprising:
an ingress module adapted to receive frames of data from a channel, wherein each
frame of data has one of a plurality of classes of service, and to store the data in one or more
5 buffers; and

an egress module adapted to exercise flow control on the channel for each of the
classes of service when the number of the buffers storing frames of data received from the
channel and having the class of service but not yet transmitted from the network switching
device exceeds a predetermined threshold for the class of service.

10

2. The network switching device of claim 1:
wherein, to exercise flow control for one of the classes of service, the egress module
is further adapted to send a pause frame to the channel, and wherein the pause frame
indicates the one of the classes of service.

15

3. The network switching device of claim 1:
wherein the egress module is further adapted to terminate flow control on the channel
for each of the classes of service when the number of the buffers storing frames of data
having the class of service but not yet transmitted from the network switching device falls
20 below a further predetermined threshold for the class of service.

25

4. The network switching device of claim 3:
wherein, to terminate flow control for one of the classes of service, the egress module
is further adapted to send a pause release frame to the channel, and wherein the pause release
frame indicates the one of the classes of service.

25

5. The network switching device of claim 1, further comprising:
one or more queues;
a forwarding module adapted to enqueue each of the buffers to one or more of the one
30 or more queues after the ingress module stores the data of one of the frames in the buffer;

a plurality of counters comprising one counter for each of the classes of service, wherein each of the counters is adapted to

store a count for the channel for a respective one of the classes of service,

increment the count when the forwarding module enqueues one of the buffers

5 storing the data from one of the frames having the respective class of service, and

decrement the count after the data stored in a buffer for a frame received from the channel and having the respective class of service is transmitted from the network switching device;

10 wherein the egress module is adapted to exercise flow control on the channel for each of the classes of service when the count for the class of service exceeds the predetermined threshold for the class of service.

6. The network switching device of claim 5:

15 wherein the egress module is further adapted to terminate flow control on the channel for each of the classes of service when the count for the class of service falls below a further predetermined threshold for the class of service.

7. An integrated circuit comprising the network switching device of claim 1.

20 8. A network switch comprising the network switching device of claim 1.

9. An output-queued network switch comprising the network switching device of claim 1.

25 10. The network switching device of claim 1, further comprising:
a memory comprising the buffers.

11. An integrated circuit comprising the network switching device of claim 10.

30 12. The network switching device of claim 1, further comprising:
a reserve module adapted to reserve one or more of the buffers to the channel;

wherein the pause threshold for the channel is a function of at least one of the group consisting of

the number of the buffers reserved to the channel; and
the number of the buffers neither reserved nor enqueued.

5

13. The network switching device of claim 3, further comprising:
a reserve module adapted to reserve one or more of the buffers to the channel;
wherein the pause release threshold for the channel is a function of at least one of the group consisting of

10 the number of the buffers reserved to the channel; and
the number of the buffers neither reserved nor enqueued.

14. A network switching device comprising:
ingress module means for receiving the frames of data from a channel, wherein each
15 frame of data has one of a plurality of classes of service, and to store the data in one or more
buffers; and
egress module means for exercising flow control on the channel for each of the
classes of service when the number of the buffers storing frames of data received from the
channel and having the class of service but not yet transmitted from the network switching
20 device exceeds a predetermined threshold for the class of service.

25

15. The network switching device of claim 14:
wherein, to exercise flow control for one of the classes of service, the egress module
means sends a pause frame to the channel; and
wherein the pause frame indicates the one of the classes of service.

30

16. The network switching device of claim 14:
wherein the egress module means terminates flow control on the channel for each of
the classes of service when the number of the buffers storing frames of data having the class
of service but not yet transmitted from the network switching device falls below a further
predetermined threshold for the class of service.

17. The network switching device of claim 16:

wherein, to terminate flow control for one of the classes of service, the egress module means sends a pause release frame to the channel, and wherein the pause release frame indicates the one of the classes of service.

5

18. The network switching device of claim 14, further comprising:

one or more queue means for queuing the buffers;

forwarding module means for enqueueing each of the buffers to one or more of the one or more queue means after the ingress module means stores the data of one of the frames in the buffer;

10

a plurality of counter means comprising one counter means for each of the classes of service, each of the counters for

storing a count for the channel for a respective one of the classes of service,

15

incrementing the count when the forwarding module enqueues one of the

buffers storing the data from one of the frames having the respective class of service, and

20

decrementing the count after the data stored in a buffer for a frame received

from the channel and having the respective class of service is transmitted from the

network switching device;

wherein the egress module is adapted to exercise flow control on the channel for each of the classes of service when the count for the class of service exceeds the predetermined threshold for the class of service.

25

19. The network switching device of claim 18:

wherein the egress module means terminates flow control on the channel for each of the classes of service when the count for the class of service falls below a further predetermined threshold for the class of service.

30

20. An integrated circuit comprising the network switching device of claim 14.

21. A network switch comprising the network switching device of claim 14.

22. An output-queued network switch comprising the network switching device of claim 14.

5

23. The network switching device of claim 14, further comprising: reserve module means for reserving one or more of the buffers to the channel; wherein the pause threshold for the channel is a function of at least one of the group consisting of

10 the number of the buffers reserved to the channel; and
the number of the buffers neither reserved nor enqueued.

24. The network switching device of claim 16, further comprising: reserve module means for reserving one or more of the buffers to the channel; wherein the pause release threshold for the channel is a function of at least one of the group consisting of

the number of the buffers reserved to the channel; and
the number of the buffers neither reserved nor enqueued.

20 25. A method comprising:

receiving frames of data from a channel, wherein each frame of data has one of a plurality of classes of service;

storing the data in one or more buffers; and
exercising flow control on the channel for each of the classes of service when the

25 number of the buffers storing frames of data received from the channel and having the class of service but not yet transmitted from the network switch exceeds a predetermined threshold for the class of service.

30 26. The method of claim 25, wherein exercising flow control for one of the classes of service comprises:

sending a pause frame to the channel, wherein the pause frame indicates the one of the classes of service.

27. The method of claim 25, further comprising:

5 terminating flow control on the channel for each of the classes of service when the number of the buffers storing frames of data having the class of service but not yet transmitted from the network switch falls below a further predetermined threshold for the class of service.

10 28. The method of claim 27, wherein terminating flow control for one of the classes of service comprises:

sending a pause release frame to the channel, and wherein the pause release frame indicates the one of the classes of service.

15 29. The method of claim 25, further comprising:

enqueueing each of the buffers to one or more output queues after storing the data of one of the frames in the buffer;

storing a count for the channel for each of the classes of service;

incrementing the count for one of the classes of service when enqueueing one of the 20 buffers storing the data from one of the frames received from the channel and having the one of the classes of service, and

decrementing the count for one of the classes of service after the data stored in a buffer for a frame received from the channel and having the one of the classes of service is transmitted;

25 wherein exercising flow control on the channel for each of the classes of service comprises exercising flow control on the channel for one of the classes of service when the count for the one of the classes of service exceeds the predetermined threshold for the one of the classes of service.

30 30. The method of claim 29, further comprising:

terminating flow control on the channel for each of the classes of service when the count for the class of service falls below a further predetermined threshold for the class of service.

5 31. A computer program embodying instructions executable by a processor to control an apparatus having an ingress module connected to a channel and an egress module connected to the channel, the computer program comprising:

 storing data in one or more buffers for frames of the data received by the ingress module from the channel, wherein each frame has one of a plurality of classes of service; and

10 causing the egress module to exercise flow control on the channel for each of the classes of service when the number of the buffers storing the data for frames received by the ingress module from the channel and having the class of service but not yet transmitted from the apparatus exceeds a predetermined threshold for the class of service.

15 32. The computer program of claim 31, wherein causing the egress module to exercise flow control for one of the classes of service comprises:

 causing the egress module to send a pause frame to the channel, wherein the pause frame indicates the one of the classes of service.

20 33. The computer program of claim 31, further comprising:

 causing the egress module to terminate flow control on the channel for each of the classes of service when the number of the buffers storing the data for frames having the class of service but not yet transmitted from the apparatus falls below a further predetermined threshold for the class of service.

25

34. The computer program of claim 33, wherein causing the egress module to terminate flow control for one of the classes of service comprises:

 causing the egress module to send a pause release frame to the channel, wherein the pause release frame indicates the one of the classes of service.

30

35. The computer program of claim 31, further comprising:

enqueueing each of the buffers to one or more output queues after storing the data for one of the frames in the buffer;

storing a count for the channel for each of the classes of service;

incrementing the count for one of the classes of service when enqueueing one of the buffers storing the data for one of the frames received from the channel and having the one of the classes of service, and

5 decrementing the count for one of the classes of service after the data stored in a buffer for a frame received from the channel and having the one of the classes of service is transmitted;

10 wherein causing the egress module to exercise flow control on the channel for each of the classes of service comprises causing the egress module to exercise flow control on the channel for one of the classes of service when the count for the one of the classes of service exceeds the predetermined threshold for the one of the classes of service.

15 36. The computer program of claim 35, further comprising:

causing the egress module to terminate flow control on the channel for each of the classes of service when the count for the class of service falls below a further predetermined threshold for the class of service.

37. A network switching device comprising:

an egress module adapted to retrieve frames of data from a memory, and to transmit the frames of data to a channel, wherein each of the frames of data has one of a plurality of classes of service; and

5 an ingress module adapted to receive a pause frame indicating one or more of the classes of service to be paused;

wherein, in response to the pause frame, the egress module is further adapted to

cease to transmit the frames of data having the one or more classes of service to be paused, and

10 continue to transmit the frames of data not having the one or more classes of service to be paused.

38. The network switching device of claim 37:

wherein the ingress module is further adapted to receive a pause release frame indicating one or more of the classes of service to be released;

15 wherein, in response to the pause frame, the egress module is further adapted to resume transmitting the frames of data having the one or more classes of service to be released.

20 39. An integrated circuit comprising the network switching device of claim 37.

40. A network switch comprising the network switching device of claim 37.

41. An output-queued network switch comprising the network switch of claim 37.

25

42. The network switching device of claim 37, further comprising:
a memory comprising the buffers.

43. An integrated circuit comprising the network switching device of claim 42.

30

44. A network switching device comprising:

egress module means for retrieving frames of data from a memory, and for transmitting the frames of data to a channel, wherein each of the frames of data has one of a plurality of classes of service; and

5 ingress module means for receiving a pause frame indicating one or more of the classes of service to be paused;

 wherein, in response to the pause frame, the egress module means

 ceases to transmit the frames of data having the one or more classes of service to be paused, and

10 continues to transmit the frames of data not having the one or more classes of service to be paused.

45. The network switching device of claim 44:

 wherein the ingress module means receives a pause release frame indicating one or more of the classes of service to be released;

15 wherein, in response to the pause frame, the egress module means resumes transmitting the frames of data having the one or more classes of service to be released.

46. An integrated circuit comprising the network switching device of claim 44.

20 47. A network switch comprising the network switching device of claim 44.

48. An output-queued network switch comprising the network switch of claim 44.

25 49. A method comprising:

 transmitting frames of data to a channel, wherein each of the frames of data has one of a plurality of classes of service;

 receiving a pause frame indicating one or more of the classes of service to be paused; and

30 in response to the pause frame,

ceasing to transmit the frames of data having the one or more classes of service to be paused, and
continuing to transmit the frames of data not having the one or more classes of service to be paused.

5

50. The method of claim 49, further comprising:
receiving a pause release frame indicating one or more of the classes of service to be released; and
in response to the pause release frame, resuming transmitting the frames of data
10 having the one or more classes of service to be released.

51. A computer program embodying instructions executable by a processor to control an apparatus having an ingress module connected to a channel and an egress module connected to the channel, the computer program comprising:
15 causing the egress module to transmit frames of data to the channel, wherein each of the frames of data has one of a plurality of classes of service; and
in response to the ingress module receiving a pause frame from the channel, the pause frame indicating one or more of the classes of service to be paused,
20 causing the egress module to cease to transmit the frames of data having the one or more classes of service to be paused, and
causing the egress module to continue to transmit the frames of data not having the one or more classes of service to be paused.

30

52. The computer program of claim 51, further comprising:
25 in response to the ingress module receiving a pause release frame from the channel, the pause release frame indicating one or more of the classes of service to be released, causing the egress module to resume transmitting the frames of data having the one or more classes of service to be released.